NASA INFORMATION SYSTEMS
SCIENCE INFORMATION SYSTEMS NEWSLETTER

Lead Feature

NASA-Inspired Parallel Work stations Exceed Sustained GigaFLOPS

Jarrett Cohen, Hughes STX, NASA Goddard Space Flight Center

Two parallel workstations based on NASA's Beowulf "pile of PCs" concept have each exceeded one gigaFLOPS performance for a price of approximately $50,000. One system was sponsored by the Office of Space Science and built by the Jet Propulsion Laboratory (JPL) and the California Institute of Technology. The other was funded and is operated by the Theoretical Division at Los Alamos National Laboratory.

Supercomputing '96 conference attendees were able to inspect the systems at the Caltech Center for Advanced Computing Research and Los Alamos research exhibits November 18–21 in Pittsburgh. Tied together with 16 additional Ethernet channels, the two Beowulfs worked in concert to realize greater than two gigaFLOPS on several demonstrations. For these achievements, Beowulf was recently nominated by the editors for the 1997 Discover Magazine Awards for Technological Innovation in the computer hardware and electronics category. Finalists will be announced in March.

About Beowulf

The newest versions of the Beowulf cluster, linking 16 Intel Pentium Pro processors with Fast Ethernet networks, attained over one gigaFLOPS sustained on a cosmological N-body simulation. The architecture was conceived by the Universities Space Research Association (USRA) Center of Excellence in Space Data and Information Sciences (CESDIS) at Goddard Space Flight Center. The NASA High Performance Computing and Communications (HPCC) Earth and Space Sciences (ESS) Project funds continuing development.

The JPL/Caltech Beowulf, which ran the N-body calculation at 1.26 gigaFLOPS, was built in collaboration with CESDIS researchers. It consists of 16 Pentium Pro (200 MHz) processors connected through a 100 megabits-per-second Fast Ethernet switch. The system has 2 gigabytes of distributed memory, a theoretical peak speed of 3.2 gigaFLOPS, and 80 gigabytes of disk storage. Michael Warren of the Theoretical Division at Los Alamos constructed a similar

machine that also relies on 16 Pentium Pro processors but contains five Fast Ethernet interfaces per processor. This system achieved 1.17 gigaFLOPS on the N-body code, which was written by Warren and Caltech's John Salmon.

The benefits

"Using commodity personal computer subsystems allows supercomputer performance at a significantly reduced cost," said Thomas Sterling, a senior scientist at Caltech and JPL who led the original design team. "Any college or university, or laboratory department, can now afford a parallel supercomputer for research and education."

In addition, "Beowulf has a larger memory and much larger disk storage than commercially available workstations in the same price range," Sterling said. "Together with processor speed these qualities provide a robust platform for applications with large datasets, such as in the Earth and space sciences."

The benchmark used by Warren and Salmon to measure the performance is a highly optimized N-body "treecode" simulating the gravitational interactions of 10 million bodies. The fastest overall implementations of the code are on 512 nodes of the Los Alamos Thinking Machines CM-5 (14.06 gigaFLOPS) and the Caltech Intel Paragon (13.70 gigaFLOPS).

Sterling summarized the significance of the Beowulf results when he stated, "We have entered a new era in which mass market commercial computing products can be harnessed for large-scale scientific computation, greatly reducing the end user cost and allowing more researchers to do more and better computational science."

Beowulf offers a sophisticated software infrastructure through an enhanced Linux operating system. Linux provides UNIX functionality on systems using Intel, Sun, and DEC Alpha processors and is widely available at no cost. CESDIS scientist Donald Becker augmented Linux with channel-bonding software to combine the performance of multiple Ethernet network channels efficiently and transparently. According to Sterling this Parallel Linux, distributed on the World Wide Web, has become the major source of Linux networking software. It also incorporates parallel programming API's such as message passing interface, parallel virtual machine, and bulk synchronous parallel.

Building and using Beowulf

NASA has instituted a Beowulf University Consortium to assist colleges and universities in building Beowulfs for teaching parallel programming techniques. Sterling pointed out that electrical and computer engineering students also can benefit from the experience of building a parallel computer. Caltech, the University of Illinois at Urbana-Champaign, and Drexel, George Washington, and Clemson Universities are current participants. Several other universities, as well as magnet high schools, have expressed interest.

Besides its use in scientific computing, Beowulf is being tested as a mass storage device and as a satellite data processing engine. A $500,000 award from the Defense Advanced Research Projects Agency is supporting a 64-node Beowulf terabyte mass storage system with a gigabyte-per-second I/O rate. This array will serve NASA HPCC's 384-node CRAY T3E being installed at Goddard this spring. Nine ESS Project Grand Challenge investigation teams will stress the system with computation runs individually producing up to 500 gigabytes of data.

NASA and the US Air Force are collaborating on placement of Beowulf workstations as inexpensive satellite readout stations. They will allow data product generation and model and forecast processing in near real time, which is a ten-fold improvement over the current standard. Planned installation sites include the NASA Regional Data Centers at Clemson, Louisiana State University, University of Hawaii, and the University of Maryland-Baltimore County.

For further information email the author or access the CESDIS Linux Beowulf homepage, respectively at:

jcohen@pop900.gsfc.nasa.gov

http://cesdis.gsfc.nasa.gov/linux-we/beowulf/beowulf.html

To view a color image of Thomas Sterling, senior scientist at Caltech and JPL, access:

http://sdcd.gsfc.nasa.gov/ESS/beopix.html

Thomas Sterling, senior scientist at Caltech and JPL, discusses the Hyglac parallel workstation with a Supercomputing ’96 attendee. Hyglac and Loki, a similar workstation at Los Alamos National Laboratory, each link 16 Pentium Pro processors to achieve gigaflops performance for $50,000. Photograph by Judy Conlon, Ames Research Center.c